

The Compactness of Construction Grammars

Wlodek Zadrozny
 IBM Research, T. J. Watson Research Center
 Yorktown Heights, NY 10598
 wlodz@watson.ibm.com *

March, 1995

Abstract

We present an argument for *construction grammars* based on the minimum description length (MDL) principle (a formal version of the Ockham Razor). The argument consists in using linguistic and computational evidence in setting up a formal model, and then applying the MDL principle to prove its superiority with respect to alternative models. We show that construction-based representations are at least an order of magnitude more compact than the corresponding lexicalized representations of the same linguistic data.

The result is significant for our understanding of the relationship between syntax and semantics, and consequently for choosing NLP architectures. For instance, whether the processing should proceed in a pipeline from syntax to semantics to pragmatics, and whether all linguistic information should be combined in a set of constraints. From a broader perspective, this paper does not only argue for a certain model of processing, but also provides a methodology for determining advantages of different approaches to NLP.

1 Introduction: Motivation and Terminology

We present an argument for a particular model for natural language, namely *construction grammars* (cf. [5] and [4], for a comprehensive introduction, and [12], [22], [23], [24], [26], [27] for computational models). The result we report establishes their optimality, in the sense of the minimum description length (MDL) principle (a formal version of the Ockham Razor). The argument consists in using linguistic and computational evidence in setting up a formal model, and then using the MDL principle to prove its superiority with respect to alternative models. The result is significant for our understanding of the relationship between syntax and semantics, and consequently for choosing NLP architectures. For instance, whether the processing should proceed in a pipeline from syntax to semantics to pragmatics, and whether all linguistic information should be combined in a set of constraints. From a broader perspective, this paper does not only argue for a certain model of processing, but also provides a methodology for determining advantages of different approaches to NLP.

After discussing the terminology, we begin our exposition with a sketch of the argument (Section 1.2). The argument will have two parts, of which the first is empirical and based on the work of other researchers on language structure and models of processing, and where the second part applies the

*IBM Research Report RC 20003 (88493)

MDL principle. Both will be presented in Section 3. For this second part we need to develop some techniques and intuitions. Thus, in Section 2, we present the details of the application of the MDL principle for a grammar of numbers. Since the domain of numbers is completely unambiguous, this exposition move will make transparent the subsequent application of MDL principle to a grammar of English in Section 3.2. We conclude the paper with a few open problems.

1.1 Terminology

By a *grammar* we mean a collection of entities, like data structures or logical formulas, that describe a formal language. Thus grammars are always formal. If the formal language in question resembles English, we call the grammar a *grammar of English* (ditto for other natural languages).

While it is possible to classify grammars along many dimensions, in this paper we are interested in two: (1) Whether a grammar is lexicalized, and (2) whether a grammar separates information about form and meaning (syntax and semantics) or mixes the two. Obviously, in each case these are mutually exclusive possibilities; i.e. either all information about language is contained in the lexicon or not, and the same for (2). If a grammar is not lexicalized, then it must refer to units larger than words, and we will call such an item of information a *structural rule*; the most typical example of which would be a phrase structure rule.

A *construction grammar* is a non-lexicalized grammar in which information about form and meaning is kept together in constructions. A *construction* is a set of constraints about form and meaning of a word, sequence of words, or a sequence of constructions. (A recursive definition).

This definition captures the idea that forms and meanings should be investigated together. However we want to say more than that. Since meaning very often depends on context, it is only natural to make this connection explicit. Therefore in the formalism we use constructions are triples:

$$< Context; Form; Meaning >$$

The *Form* describes the construction as a combination of subconstructions (e.g. a noun phrase as a combination of a numeral and a noun); the *Meaning* part specifies how the meanings of the subparts contribute to the meaning of the construction (e.g. specifying the number of elements). The *Context* specifies the parameters that are necessary to construct the meaning, and which are not present in the meanings of the parts; for example, the content of the question is necessary to construct the meaning of an answer, especially if the answer is just a sentence fragment.

The *minimum description length (MDL) principle* was proposed by Rissanen [19]. It states that the best theory to explain a set of data is the one which minimizes the sum of

- the length, in bits, of the description of the theory, and
- the length, in bits, of data when encoded with the help of the theory.

In our case, the data is the language we want to describe, and the the encoding theory is its grammar (which includes the lexicon). The MDL principle justifies the intuition that a more compact grammatical description is better. At issue is what is the best encoding. To address it, we will be simply comparing two complementary classes of encodings and showing that one of them is usually more compact. The formal side of the argument will be kept to the minimum: after building the two complementary models of language, the mathematics will be simple — counting. (But we will discuss some ways of producing more refined models).

1.2 The line of the argument

The paper is about how to best represent information about language, i.e. about data structures. As we know, data structures are determined by both the types of data and their structure. Therefore we have to argue for particular types and particular structures. In our particular case, to argue for grammars of constructions, we will show that the data types should contain information about both form and meaning. Secondly, we will show that their structure should contain something resembling phrase structure rules; we do it by presenting an MDL-based argument against lexicalized representation of forms and meanings.

The optimality argument goes as follows: we have to prove that for NLP it is preferable not to separate syntactic, semantic and pragmatic information — which is an argument for having data structures that combine them. But we could imagine say lexicalized grammars in which such information is combined. Hence, as the second step, we have to show that grammars with "phrase structure rules" are better than lexicalized grammars. These two arguments show the superiority of a construction-based approach with respect to alternative grammatical formalism.

2 Using MDL with grammars of numbers

To make the presentation clear we first discuss a grammar of numbers — for numbers our whole argument is completely formal and transparent. In the next section we use the same argument for NL grammars.

A lexicalized grammar of numbers

Any grammar of numbers must somehow express the fact that the value of a digit D depends on its position. I.e.

$$\mu(D) = D * 10^{*pos(D)}$$

where $pos(D)$ is the position of D (counting from the right, beginning with 0). Notice that a lexicalized grammar expresses it directly, and must repeat it for every digit. (We use the same symbols for the digits and their values).

$$< [10]; 0; \mu(0) = 0 >$$

$$< [10]; 1; \mu(1) = 1 * 10^{*pos(1)} >$$

$$< [10]; 2; \mu(2) = 2 * 10^{*pos(2)} >$$

...

$$< [10]; 9; \mu(9) = 9 * 10^{*pos(9)} >$$

Based on that, the value of a number is the sum of values of its digits:

$$\mu(N) = \sum_{D \text{ in } N} \mu(D)$$

Note that the grammar specifies the formula for value of the type, and the value of the token is given by its instantiation. E.g. in computing $\mu(17341)$, notice the different values of the digit 1. Also, note that given a different set of functions, e.g. *head*, *tail*, *log*, ***, *+*, the μ function would be slightly different, but the lexicon must look essentially the same, because there is no other place to

put the data about how the forms determine the meanings.

A construction grammar of numbers

The lexical part of a construction grammar can now be much simpler. (And it could be simplified even further by assuming that the value of a token is the token, unless specified otherwise).

$$< [10]; 0; \mu(0) = 0 >$$

$$< [10]; 1; \mu(1) = 1 >$$

$$< [10]; 2; \mu(2) = 2 >$$

...

$$< [10]; 9; \mu(9) = 9 >$$

In contrast to the lexicalized grammar, in a grammar of constructions we can write a structural rule

$$< [10]; DS \rightarrow DS1 D ; \mu(DS) = 10 * \mu(DS1) + \mu(D) >$$

This rule defines a production saying that a new structure is obtained by adding a digit D to a previously defined structure $DS1$. The equality associates the meaning of a new structure $\mu(DS)$ with the meaning of its components $\mu(DS1)$ and $\mu(D)$. As a consequence of choosing this kind of representation, the rule about how to compute the meaning of digits has to be stated only once.

As we can see, describing the same language with a grammar of constructions results in a more compact grammar. We saved 11 symbols per non-0 lexical entry, i.e. 99 symbols altogether. Although, we added a structural rule, its size is comparable with the above Σ rule for computing the value of a sequence of digits in the lexicalized grammar. Also note that the latter must additionally refer to the function *pos* and exponentiation.

While savings 99 symbols is not much, for larger lexicons the saving would be much bigger. Larger lexicons are obtained by increasing the *Base*. In this case the grammar production (phrase structure rule) reads:

$$< [Base]; DS \rightarrow DS1 D ; \mu(DS) = Base * \mu(DS1) + \mu(D) >$$

It can be easily checked that the resulting construction grammar is always an order of magnitude more compact than its lexicalized counterpart.

A grammar of constructions is even more compact if additional conditions are placed on sequences of digits, e.g. that only ascending sequences of digits are acceptable. That is so, because typically any such a condition would have to be included in all lexical entries, and the more complicated the condition, the less compact is the lexicalized grammar.

3 The superiority of constructions for NLU grammars

In this section we present the argument that (a) it makes sense to encode syntactic, semantic and pragmatic information together, and (b) that construction-based grammar are more compact than lexicalized grammar that encode the same semantic information. At the end of the section we discuss some possible extensions of the model to cover the case of "lexicalized grammars with a few constructions".

3.1 Data types: FORMS, MEANINGS, and CONTEXTS

There are three arguments supporting the encoding of linguistic information in data structures that combine syntactic, semantic and pragmatic information. Namely, the linguistic theory, the practice of computational linguists who encode it that way, and experimental evidence from analyzing parsing mechanisms.

We now briefly discuss each argument. Regarding linguistic theory, the *Comprehensive Grammar of the English Language* [20] describes the language by freely combining syntax, semantics and pragmatics. With a closer look, the analysis of the structure of VPs and NPs requires the reference to semantic information; e.g. McCawley [15], vol.1, p.222 and ff. argues that the restriction on the progressive *be* is semantic rather than syntactic in nature, that is, its complement should refer to an **activity** or a **process** rather than a **state**. Finally, [5] contains a comprehensive set of arguments for construction-based description of clause structure.

Moving to relevant work in computational linguistics, first let us notice that in some cases the strategy of describing language structure using only syntactic markers produces impressive results (e.g. [21]). But when the goal is to understand language, with the increasing coverage of a grammar the set of its markers grows and encodes more and more semantics. For instance PEG, a broad coverage grammar of English, [9], [10], used about 400 markers, including MONEY, DATE, PHONE, ANIMATE, HUMAN, RELIGIOUS_NAME, TIME, TITLE, VERB_OF_COGNITION, VERB_OF_ASKING, DEROGATORY, EMPHATIC etc. Furthermore, if we examine computational linguistic literature we can see that to prevent overgeneralizations, semantic and pragmatic information must be taken into account (see e.g. [1], [6], [7], [11]); and it is virtually impossible to produce a correct predicate-argument structure for many constructions (e.g. PPs; relative clauses; parallel, cumulative or periodic sentences) without incorporating those two kinds of information (see also [3], [18], [17], [13]). Clearly, the integration of various types of information is necessary to interpret discourse (e.g. [8]).

While the body of work we have quoted provides evidence from coverage and depth of semantic processing, there is also evidence based on the efficiency and robustness of parsing. Thus, Lytinen [14] shows that a semantically driven approach is superior to syntax-first approach in processing text in narrow domains. His semantic-first algorithm decides which grammar rules to apply next on the basis of "desirable semantic attachments between adjacent constituents", after such constituents have already been identified. Dowding et al. [2] show that the same strategy works with a large coverage grammar of English.

3.2 The structure of data: Why the argument for numbers works also for NL

To show that the argument of Section 2 works for a grammar of English we should show that the use of structural rules can result in a more compact grammar.

We need a simple model, so let us consider the problem of determining whether a clause followed by a set of PPs is nonsensical. For example,

we meet at 12 with bob at 6 avenue and 44 street vs.

the dow closed at 2200 with bob at 6 avenue and 44 street.

In general, this problem cannot be solved without access to a large body of background knowledge; so, let us simplify it further. Assume that sentences that repeat the same type of information are nonsensical, e.g.

we meet at 12 pm with bob from 5 to 6 pm

To set up the model we have to define a formal language resembling English. We do it in two steps. Let our first formal language L_1 consist of all sequences of SV (subject-verb), SVO (subject-verb-object), SVOO (subject-verb-object-indirect object), of English taken from some very large corpus.

The language L_{PP} we are interested in consists of sentences of L_1 followed by any number of PPs (prepositional phrases) that contain the nouns from L_1 . (Hence L_{PP} is infinite). Using this model we can discuss differences between a construction grammar and a lexicalized grammar.

Construction grammar

To define a construction grammar for L_{PP} we define the lexicon and a set of productions (phrase structure rules). Let the lexical entries be given by the matrix:

$$< \square; w; \mu(w) = \{cat(type_i, w_i) : i < n_w\} >$$

i.e. the meaning of a word is given by its linguistic category, a word sense, and the semantic type of the word sense. The meaning is given as a set, because words often belong to different categories and may have multiple word senses, e.g.

$$\mu(dog) = \{verb(pursue, dog_0), noun(person, dog_1), noun(person, dog_2), noun(animal, dog_3)\}$$

Remark. We ignore the context to make the argument more general. Also, we could write $< \square; w; \mu(w) = cat(w_i) >$, and associate one word with multiple lexicon entries (slightly abusing the notation).

Having defined the lexicon, we have to cover the SV, SVO, and SVOO constructions; for instance we could write the SVO-action construction as

$$< \square; CL \rightarrow S V O ; \mu(CL) = [[action, \mu(V)], [agent, \mu(S)], [object, \mu(O)]] >$$

However for our purposes it is irrelevant how the meanings of those constructions are encoded; the only thing we will need is the existence of the μ function.

The next step is to cover the PPs. We represent them as

$$< \square; pp \rightarrow prep(p) noun(type_i, w_i); \mu(pp) = pp(type_{i,p}, w_i) >$$

I.e. we assume that prepositional phrases have types, and each type is a function of the preposition and the noun type. E.g.

$$< \square; pp \rightarrow prep(at) noun(hour, X); \mu(pp) = pp(event_time, X) >$$

Thus "at" followed by the second sense of 2200 (as hour, not number), would be classified as an "event_time". (We are conveniently assuming that numerals are nouns).

Finally, we define the construction of adding an adjunct to a clause (with + standing for the simple *append* or a more complex procedure):

$$< \square; CL \rightarrow CL1 A ; \mu(CL) = \mu(CL1) + \mu(A) \quad \text{if } \mu(A) \text{ not in } \mu(CL1), \text{ and } \perp - \text{otherwise} >$$

Thus defined construction grammar encodes our formal language L_{PP} .

Lexicalized grammar

For a lexicalized grammar, as in Section 2, we observe that the meaning of a clause such as the one above must be encoded with each noun. Thus, for each noun we have to encode the meaning of its adjuncts, i.e. repeat the formula:

"if I am combined with preposition X_1 , then our type will be $T = T_1$;
 "if I am combined with preposition X_2 , then our type will be $T = T_2$;
 ...
 and if T does not appear in the meaning E of everything to the left
 of the preposition immediately to the left of me, then I make sense;
 and our joint meaning of is the sum of T and E ."

For instance, for each number that can denote an hour we would have:

$$\begin{aligned}
 < [] ; 2 ; \quad \mu(2) = \{(hour, 2)\} \\
 &\quad \text{if "at" then } T = event_time \\
 &\quad \text{if "from" then } T = beginning_time \\
 &\quad \dots \\
 &\quad \text{if } T \text{ appears in } E \text{ to the left then } \perp \\
 &\quad \text{otherwise add } T \text{ to } E >
 \end{aligned}$$

(In addition the formulas might encode constraints on *event.time*, *beginning.time*, and other types). As before, the reason for repeating this piece of information is that these formulas must somehow be encoded. Since the reference to anything larger than a word is forbidden, the formula must be repeated for every word separately.

Repeating the argument of Section 2, we conclude that a construction grammar that encodes the formal language L_{PP} is at least an order of magnitude more compact than any lexicalized grammar that encodes this language. The exact difference in bits can be computed as in [16] pp. 312-316, as a function of the number of nouns in the language.

3.3 Why adding a few structure rules won't help

With any mathematical model of language one should ask the question how closely it approximates the linguistic reality, and what happens if we introduce some minor changes to it. For example, what if a grammar is 99.99% lexicalized and contains only a few phrase structure rules (e.g. the rule about the meanings of clauses with adjuncts). The answer is that the same argument would apply, provided we can find another productive grammatical phenomenon that behaves similarly to digits in our grammar of numbers or to PPs in the case of L_{PP} . Is there a reason to believe that such productive phenomena are common? Yes.

The term to denote a situation where a language pattern is productive, but has many exceptions, is *partial productivity* ([5]). One such pattern is the ditransitive construction (*ibid.*):

John faxed Bob the report.

Barbara told Bob the story.

**Barbara whispered Bob the story.*

To account for its partial productivity, a rule that describes its usage must refer to the semantic types of the verb, the object, and the indirect object. This situation is similar to the one described in Section 3.2. Thus, the ditransitive construction should not be lexicalized either.

Most linguistic phenomena exhibit this quality of partial productivity. And to many the techniques of Section 3.2 could be applied. These are for instance adjuncts, open idioms [4], noun-noun modifications, the order of adjectives, and long distance dependencies [12]. For all of them we could set up models similar to the ones discussed above; furthermore, these models could be combined in a bigger structure. Thus, by using a more comprehensive model of English, we could have

made the construction grammars many orders of magnitude more compact than their lexicalized counterparts.

This argument can be even further strengthened by looking beyond the clause. Similar models can be created for connectives and linking words (such as "therefore", "consequently"). Moving to dialogs, where understanding sentence fragments requires the context (e.g. about previous discourse), we could build models, where the difference in the size of construction grammars and a lexicalized grammars would be due to the necessity to encode the context (the simple task being e.g. to identify reasonable and absurd answers). Using the techniques of [25] it is easy to show that the arguments of Sections 2 and 3 could also be used with paragraph coherence, and most likely extended beyond the paragraph.

4 Open problems and Conclusions

Since this paper is about counting, the most natural open problem is the number of constructions beyond the word, and then beyond the clause. Of course, it is not clear how to count them. Still, we know approximately the size of lexicons for spellchecking; similarly, we could ask about the size of a construction grammar for NLU (e.g. at the level of achieving a decent score on a SAT exam) or for summarizing of New York Times stories.

The next question is whether using techniques similar to those of Section 3.2 it can be shown that some particular ways of representing constructions are optimal (e.g. the embedding of ontology into the grammar or the restriction to only two levels of structure, as in [24]).

Finally, there are many questions that are not directly related to the topic of this paper, e.g. the role of ontologies in describing constructions, the interaction between structural and functional descriptions etc.

Conclusions: We used a combination of empirical data and the minimum description length principle to show that construction grammars are better representation of linguistic data than their lexicalized counterparts. This hybrid — linguistic, computational, and mathematical — argument for a construction-based approach to language understanding is the main contribution of this paper. However, it is likely similar arguments could be used in other circumstances where one has to choose between competing representations.

References

- [1] J. Allen. *Natural Language Understanding*. Benjamin/Cummings, Menlo Park, CA, 1987.
- [2] J. Dowding, R. Moore, F. Andry, and D. Moran. Interleaving syntax and semantics in an efficient bottom-up parser. In *Proc. of 32nd Annual Meeting of the Association for Computational Linguistics, ACL*. ACL, 1994.
- [3] J.E. Fenstad, P-K. Halvorsen, T. Langholm, and J. van Benthem. *Situations, Language and Logic*. D. Reidel, Dordrecht, Holland, 1987.
- [4] C.J. Fillmore, P. Kay, and M.C. O'Connor. Regularity and idiomaticity in grammatical constructions. *Language*, 64(3):501–538, 1988.
- [5] A.E. Goldberg. *Constructions: a construction grammar approach to argument structure*. The University of Chicago Press, Chicago, IL, 1994.

- [6] E. A. Hinkelman and J.F. Allen. Two constraints on speech acts ambiguity. In *Proc. of 27th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 212–219, Vancouver, BC, 1989. ACL.
- [7] G. Hirst. *Semantic interpretation and the resolution of ambiguity*. Cambridge University Press, Cambridge, Great Britain, 1987.
- [8] J. R. Hobbs, M.E. Stickel, D.E. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63(1-2):69–142, 1993.
- [9] K. Jensen. Parsing strategies in a broad-coverage grammar of english. Technical Report RC 12147, IBM T.J. Watson Research Center, Yorktown Heights, New York, 1986.
- [10] K. Jensen, G.Heidorn, and S.D.Richardson, editors. *Natural Language Processing: The PLNLP Approach*. Kluwer, Boston, 1993.
- [11] K. Jensen and J.-L. Binot. Disambiguating prepositional phrase attachments by using on-line dictionary definitions. *Computational Linguistics*, 13(3-4):251–260, 1988.
- [12] D. Jurafsky. *An On-line Computational Model of Sentence Interpretation*. PhD thesis, University of California, Berkeley, 1992. Report No. UCB/CSD 92/676.
- [13] E.L. Keenan. Relative clauses. In T. Shopen, editor, *Language Typology and syntactic description Vol.II.*, pages 140–170. Cambridge University Press, New York, NY, 1985.
- [14] S.L. Lytinen. Semantics-first natural language processing. In *Proceedings AAAI-91*, pages 111–116, Anaheim, CA, 1991.
- [15] J.D. McCawley. *The Syntactic Phenomena of English*. University of Chicago Press, Chicago, IL, 1988.
- [16] M.Li. and P.Vitanyi, editors. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 1993.
- [17] M. Noonan. Complementation. In T. Shopen, editor, *Language Typology and syntactic description Vol.II.*, pages 42–140. Cambridge University Press, New York, NY, 1985.
- [18] Y. Ravin. *A Decompositional Approach to Predicates Denoting Events*. PhD thesis, CUNY, New York, 1987.
- [19] J. Rissanen. A universal prior for integers and estimation by minimum description length. *Annals of Statistics*, 11:416–431, 1982.
- [20] R.Quirk, S.Greenbaum, G.Leech, and J.Svartvik. *A Comprehensive Grammar of the English Language*. Longman, New York, NY, 1985.
- [21] P. Tapainen and A. Voutilainen. Tagging accurately – don’t guess if you know. In *Proc. Fourth Conference on Applied Natural Language Processing*, pages 47–52, Stuttgart, Germany, 1994.
- [22] N. Ward. A parallel approach to syntax for generation. *Artificial Intelligence*, 57(2-3):183–225, 1992.
- [23] W. Zadrozny. From utterances to situations: Parsing with constructions in small domains. In *Language, Logic and Computation: The 1994 Moraga Proceedings*. CSLI, Stanford, 1994.

- [24] W. Zadrozny. Parsing with constructions and ontology. In *Proc. AAAI Fall Symposium on Knowledge Representation for Natural Language Processing in Implemented Systems*. AAAI, New Orleans, Louisiana, 1994.
- [25] W. Zadrozny and K. Jensen. Semantics of paragraphs. *Computational Linguistics*, 17(2):171–210, 1991.
- [26] W. Zadrozny and A. Manaster-Ramer. The significance of constructions. Technical Report RC 20002 (88492), IBM Research, T.J. Watson Research Center, 1995. (a 1993 manuscript).
- [27] W. Zadrozny, M. Szummer, S. Jarecki, D.E. Johnson, and L. Morgenstern. NL understanding with a grammar of constructions. *Proc. Coling’94*, 1994.